Working on the Jean Zay machine

Table of Content

Last Update 10/10/2019

1. Introduction

- On-line users manual: ■http://www.idris.fr/eng/jean-zay
- Jean-Zay computing nodes: the nodes of CPU partition have 40 cores each.
  - Intel Cascade Lake nodes for regular computation
  - Partition name: **cpu_p1**
  - CPUs: 2x20-cores Intel Cascade Lake 6248 @2.5GHz
  - Cores/Node: 40
  - Nodes: 1 528
  - Total cores: 61120
  - RAM/Node: 192GB
  - RAM/Core: 4.8GB
- Jean-Zay post-processing nodes : xlarge are free and useful for post-processing operations.
  - Fat nodes for computation requiring a lot of shared memory
  - Partition name: **prepost**
  - CPUs: 4x12-cores Intel Skylake 6132@3.2GHz
  - GPUs: 1x Nvidia V100
  - Cores/Node: 48
  - Nodes: 4
  - Total cores: 192
  - RAM/Node: 3TB
  - RAM/Core: 15.6GB

2. Job manager commands

- `sbatch job` -> submit a job
- `scancel ID` -> kill the job with the specified ID number
- `sacct -u login -S YYYY-MM-DD` -> display all jobs submitted by login, add `-f` to see full job name
- `squeue` -> display all jobs submitted on the machine.
- `squeue -u $(whoami)` ->display your jobs.

3. Example of a job to start an executable in a Parallel environnement

3.1. MPI

Here is an example of a simple job to start an executable orchidee_ol (or gcm.e commented). The input files and the executable must be in the directory before starting the executable.

```
#!/bin/bash
#SBATCH --job-name=TravailMPI      # name of job
#SBATCH --ntasks=80                # total number of MPI processes
#SBATCH --ntasks-per-node=40       # number of MPI processes per node
# /!\ Caution, "multithread" in Slurm vocabulary refers to hyperthreading.
#SBATCH --hint=nomultithread       # 1 MPI process per physical core (no hyperthreading)
#SBATCH --time=00:10:00            # maximum execution time requested (HH:MM:SS)
#SBATCH --output=TravailMPI%j.out  # name of output file
#SBATCH --error=TravailMPI%j.out   # name of error file (here, in common with output)

# go into the submission directory
cd ${SLURM_SUBMIT_DIR}


# echo of launched commands
```

```
set -x

# code execution
srun ./orchidee_ol
#srun ./gcm.e
```

3.2. Hybrid MPI-OMP

```
#!/bin/bash
#SBATCH --job-name=Hybrid          # name of job
#SBATCH --ntasks=8             # name of the MPI process
#SBATCH --cpus-per-task=10     # number of OpenMP threads
# /!\ Caution, "multithread" in Slurm vocabulary refers to hyperthreading.
#SBATCH --hint=nomultithread   # 1 thread per physical core (no hyperthreading)
#SBATCH --time=00:10:00            # maximum execution time requested (HH:MM:SS)
#SBATCH --output=Hybride%j.out     # name of output file
#SBATCH --error=Hybride%j.out      # name of error file (here, common with the output file)

# go into the submission directory
cd ${SLURM_SUBMIT_DIR}


# echo of launched commands
set -x

# number of OpenMP threads
export OMP_NUM_THREADS=$SLURM_CPUS_PER_TASK
# OpenMP binding
export OMP_PLACES=cores

# code execution
srun ./lmdz.e
```

3.3. MPMD